Rec'd PCT/PTO 15 MAR 2005
05. 08. 03

#2

Europäisches Patentamt | European Patent Office | Office européen des brevets

REC'D 02 SEP 2003

Bescheinigung | Certificate | Attestation

Die angehefteten Unterlagen stimmen mit der ursprünglich eingereichten Fassung der auf dem nächsten Blatt bezeichneten europäischen Patentanmeldung überein.

The attached documents are exact copies of the European patent application described on the following page, as originally filed.

Les documents fixés à cette attestation sont conformes à la version initialement déposée de la demande de brevet européen spécifiée à la page suivante.

Patentanmeldung Nr. Patent application No. Demande de brevet n°

02078848.5

PRIORITY DOCUMENT
SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH
RULE 17.1(a) OR (b)

Der Präsident des Europäischen Patentamts;
Im Auftrag

For the President of the European Patent Office

Le Président de l'Office européen des brevets
p.o.

R C van Dijk

Anmeldung Nr:
Application no.: 02078848.5
Demande no:

Anmeldetag:
Date of filing: 17.09.02
Date de dépôt:

Anmelder/Applicant(s)/Demandeur(s):

Koninklijke Philips Electronics N.V.
Groenewoudseweg 1
5621 BA Eindhoven
PAYS-BAS

Bezeichnung der Erfindung/Title of the invention/Titre de l'invention:
(Falls die Bezeichnung der Erfindung nicht angegeben ist, siehe Beschreibung.
If no title is shown please refer to the description.
Si aucun titre n'est indiqué se referer à la description.)

A method of synthesis for a steady sound signal

In Anspruch genommene Priorität(en) / Priority(ies) claimed /Priorité(s) revendiquée(s)
Staat/Tag/Aktenzeichen/State/Date/File no./Pays/Date/Numéro de dépôt:

Internationale Patentklassifikation/International Patent Classification/
Classification internationale des brevets:

G10L13/00

Am Anmeldetag benannte Vertragstaaten/Contracting states designated at date of filing/Etats contractants désignées lors du dépôt:

AT BE BG CH CY CZ DE DK EE ES FI FR GB GR IE IT LI LU MC NL PT SE SK TR

A method of synthesis for a steady sound signal

The present invention relates to the field of synthesizing of speech or music, and more particularly without limitation, to the field of text-to-speech synthesis.

The function of a text-to-speech (TTS) synthesis system is to synthesize speech from a generic text in a given language. Nowadays, TTS systems have been put into practical operation for many applications, such as access to databases through the telephone network or aid to handicapped people. One method to synthesize speech is by concatenating elements of a recorded set of subunits of speech such as demisyllables or polyphones. The majority of successful commercial systems employ the concatenation of polyphones. The polyphones comprise groups of two (diphones), three (triphones) or more phones and may be determined from nonsense words, by segmenting the desired grouping of phones at stable spectral regions. In a concatenation based synthesis, the conversation of the transition between two adjacent phones is crucial to assure the quality of the synthesized speech. With the choice of polyphones as the basic subunits, the transition between two adjacent phones is preserved in the recorded subunits, and the concatenation is carried out between similar phones.

Before the synthesis, however, the phones must have their duration and pitch modified in order to fulfil the prosodic constraints of the new words containing those phones. This processing is necessary to avoid the production of a monotonous sounding synthesized speech. In a TTS system, a prosodic module performs this function. To allow the duration and pitch modifications in the recorded subunits, many concatenation based TTS systems employ the time-domain pitch-synchronous overlap-add (TD-PSOLA) (E. Moulines and F. Charpentier, "Pitch synchronous waveform processing techniques for text-to-speech synthesis using diphones," Speech Commun., vol. 9, pp. 453-467, 1990) model of synthesis. When the signal to be synthesized is required to have an extended duration this is accomplished by repeating the pitch bells, which have been obtained from the original signal. This repetition process is illustrated in Fig. 1. Time axis 100 belongs to the time domain of the original signal. The original signal has a length of T spanning the time interval between zero and T on the time axis 100. Further, the original signal has a fundamental frequency f, which corresponds to a period p; pitch bells are obtained from the original signal by

windowing the original signal by means of windows 102. In the example considered here the windows are spaced apart by the period p in the domain of time axis 100. This way the pitch bell locations i are determined on time axis 100. Time axis 104 belongs to the time domain of the signal to be synthesized. The signal to be synthesized is required to have a duration of yT, where y can be any number. Next a number of pitch bell locations j is determined on the time axis 104. Like on the time axis 100, the pitch bell locations j are spaced apart by the period p corresponding to the fundamental frequency f of the original signal. In order to increase the duration of the original signal each of the original pitch bells obtained from the original signal is repeated a number of y.times. This results in a number of intervals 106, 108, ... in the domain of time axis 104, whereby each of the intervals 106, 108, ... is composed of repetitions of identical pitch bells. For example the interval 106 contains repetitions of the pitch bell obtained from the pitch bell location i = 1 from the original signal at pitch bell locations j (i = 1, k = 1) to j (i = 1, k = y). This means that interval 106 contains a number of y repetitions of the pitch bell obtained from pitch bell location i=1 on time axis 100 of the original signal. Likewise the following interval 108 contains a number of y repetitions of the pitch bell obtained from pitch bell location i=2 from the original signal. As a consequence the synthesized signal is composed of concatenated sequences of pitch bell repetitions.

A common disadvantage of such PSOLA methods is that an extreme duration manipulation introduces audible transitions between the sequences into the signal. In particular this is a problem when the original sound is a hybrid sound like voiced fricatives having both a noisy and a periodic component. The repetition of pitch bells introduces periodicity in the noisy components, which makes the synthesized signal sound unnatural.

The present invention therefore aims to provide an improved method of synthesizing a sound signal, in particular for extreme duration modifications, like for singing.

The present invention provides for a method of synthesizing a sound signal based on an original signal in order to manipulate the duration of the original signal. In particular, the present invention enables extreme duration and pitch modifications of the original signal without audible artefacts. This is especially useful for synthesizing of singing where extreme duration manipulations in the order of 4 to 100 times of the original signal can occur.

In essence, the present invention is based on the observation that prior art PSOLA methods introduce artefacts into a synthesized signal after duration manipulation because the transition from one chain of repeating pitch bells to the next is audible. This effect which is experienced when a prior art PSOLA type method is employed for extreme

duration manipulations is particularly detrimental for hybrid sounds containing both a noisy and a periodic component.

In accordance with the invention, pitch bells are randomly selected from the original signal for each of the required pitch bell locations of the signal to be synthesized. This way the introduction of periodicity in the noisy components can be avoided and the naturalness of the original sound is preserved. In accordance with a preferred embodiment of the invention the original sound is a voiced fricative having both a noisy and a periodic component. Application of the present invention to such voiced fricatives is especially beneficial.

In accordance with a further preferred embodiment of the invention a raised cosine is used for windowing of voiced fricatives. For unvoiced sound intervals a sine window is used which has the advantage that the total signal envelope in power domain remains about constant. Unlike a periodic signal, when two noise samples are added, the total sum can be smaller than the absolute value of any of the two samples. This is because the signals are (mostly) not in-phase; the sine window adjusts for this effect and removes the envelope-modulation.

In accordance with a further preferred embodiment of the invention the original sound signal has periods which are spectrally alike and which have basically the same information content. Such periods, which are voiced, are classified by a first classifier and such periods which are unvoiced are classified by means of a second classifier.

In accordance with a further preferred embodiment of the invention the classification information of the original signal is stored in a computer system, such as a text-to-speech system. Intervals of the original signal which are classified as voiced or unvoiced steady periods being spectrally alike are processed in accordance with the present invention whereby a raised cosine window is used for voiced intervals and a sine window is used for unvoiced intervals.

In the following preferred embodiments of the invention are described in greater detail by making reference to the drawings in which:

Fig. 1 is illustrative of a prior art PSOLA-type method,

Fig. 2 is illustrative of an example for synthesizing a sound signal in accordance with an embodiment of the present invention,

Fig. 3 is illustrative of a flow chart of an embodiment of a method of the present invention,

Fig. 4 shows an example of an original signal and of the synthesized signal, and

Fig. 5 is a block diagram of a preferred embodiment of a computer system

Fig. 2 shows an example of synthesizing a signal based on an original signal. Time axis 200 is illustrative of the time domain of the original signal. The original signal has a duration T and spans the time between zero and T on time axis 200. The original signal has a fundamental frequency f which corresponds to a period p. The period p determines locations i on time axis 200 for windowing of the original signal by means of window 202. In the example considered here, the original signal is a voiced hybrid sound such that a cosine window in accordance with the following formula is used.

$$w[n] = 0.5 - 0.5 \cdot \cos\left(\frac{2\pi \cdot (n+0.5)}{m}\right), \qquad 0 \le n < m$$

In previous relation, m is the length of the window and n is the running index.

When the original signal is an unvoiced sound signal it is preferred to use the following window.

$$w[n] = \sin\left(\frac{\pi \cdot (n+0.5)}{m}\right), \qquad 0 \le n < m$$

The time domain of the signal to be synthesized is illustrated by time axis 204. The signal to be synthesized is required to have a duration of yT, where y can be any number, for example y= 4 or y = 6 or y = 20 or y = 50 or y = 100.

The period p does also determine the pitch bell locations j on time axis 204. Like on time axis 200 the pitch bell locations are spaced apart by period p. For each of the required pitch bell locations j, a random selection of a location of a pitch bell i in the time domain of the time axis 200 is made. In the example considered here there is a number of 6 pitch bells which are obtained by windowing of the original signal in the time domain of time axis 200. To select one of these obtained pitch bells for a pitch bell location j a random number between 1 and 6 is generated. This way a random selection from the available pitch bells on pitch bell locations i = 1 to i = 6 is made. This process is repeated for all required pitch bell locations j on time axis 204. For example a pitch bell for the required pitch bell

location j = 1 is selected by generating a random number between 1 and 6. In the example considered here, the number 6 is obtained such that the pitch bell obtained from pitch bell location i = 6 on the time axis 200 is selected for the required pitch bell location j = 1 on the time axis 204. Likewise a random number is generated for the required pitch bell location j = 2. The random number is 4 in this example such that the pitch bell at pitch bell location i = 4 on time axis 200 is selected for the required pitch bell location j = 2. This process is performed for all required pitch bell locations j = 1 to j = z on time axis 204. Due to the random selection of the pitch bells from the domain of the original signal, intervals 106, 108, . . . are avoided (cf. Fig. 1). As a consequence no such artefact is introduced into the synthesized signal and the synthesized signal sounds naturally even for extreme duration manipulations.

Fig. 3 shows a flow chart, which is illustrative of this method. In step 300 a recording of an original sound is provided. In step 302 hybrid sound intervals are identified and classified as voiced or unvoiced in the original sound recording. This can be done manually by a human expert or by means of a computer program, which analyses the original signal and/or its frequency spectrum for steady periods. Preferably the first analysis is performed by means of a program and a human expert reviews the output of a program. In step 304 pitch bells are obtained from the original sound signal by means of windowing. Windowing is performed by means of windows which are positioned synchronously with the fundamental frequency of the original sound signal, i.e. the windows are distanced by the period p of the original sound signal in the domain of the original sound signal. In step 306 the pitch bell locations j for which pitch bells are required in order to synthesize the signal are determined. Again the required pitch bell locations j are distanced by the period p. Alternatively the pitch bell locations j can be distanced by another period q corresponding to a higher or lower required fundamental frequency of the signal to be synthesized. This way the duration and the frequency can be modified. In step 308 a random selection of pitch bells is made for each of the required pitch bell locations j within the sound interval which is classified as hybrid. For other sound intervals a prior art PSOLA-type method may or may not be employed. In step 310 the pitch bells are overlapped and added on the pitch bell locations j in the domain of the signal to be synthesized.

Fig. 4 shows an example of an original sound signal 400 which is a diphone of /z/ to /z/transition. Also the frequency spectrum 402 of the sound signal 400 is shown in Fig. 4.

Sound signal 404 is obtained from sound signal 400 in accordance with the present invention by randomly selecting pitch bells obtained from the sound signal 400 for the required pitch bell locations in the time domain of the synthesized sound signal 404. In the example considered here the synthesized sound signal 404 is y=5 times longer than the original sound signal 400. Also the frequency spectrum 406 of the sound signal 404 is shown in figure 4. As apparent from the sound signal 404 and its frequency spectrum 406 the characteristics of the original sound signal 400 are preserved in the synthesized signal and no artefacts are introduced. As a consequence the sound signal 404 sounds identical to the sound signal 400 but is 5 times longer.

Fig. 5 shows a block diagram of a computer system, such as a text-to-speech synthesis system. The computer system 500 comprises a module 502 for storing of an original sound signal. Module 504 serves to enter and store sound classification information for the original sound signal stored in module 502. For example, steady voiced periods are marked with an 'r' and steady unvoiced periods are marked with an 's' in the original sound signal. Module 506 serves for windowing of the original sound signal of module 502 in order to obtain pitch bells. Depending on the sound classification a raised cosine or a sine window is used for steady voiced periods or steady unvoiced periods, respectively. Module 508 serves to determine the required pitch bell locations j in the time domain of the signal to be synthesized. In order to determine the required pitch bell locations j the input parameter 'length y' is utilized. The input parameter length y specifies the multiplication factor for the duration of the original signal. Further it is possible to provide a dynamically varying pitch as an additional input parameter to modify the fundamental frequency in addition to or instead of the duration.

Module 510 serves to select pitch bells from the set of pitch bells obtained from the original sound signal. Module 510 is coupled to pseudo random number generator 512. For each of the required pitch bell locations in the domain of the signal to be synthesized, a pseudo random number is generated by pseudo random number generator 512. By means of these random numbers selections of pitch bells from the set of pitch bells are made by module 510 in order to provide a randomly selected pitch bell for each of the required pitch bell locations in the time domain of the signal to be synthesized. Module 514 serves to perform an overlap and add operation on the selected pitch bells in the time domain of the signal to be synthesized. This way the synthesized signal having the required duration is obtained.

It is to be noted that the present invention can be applied on steady regions. For example, such a steady region can be a vowel or a noisy voiced sound like /z/. Hence, the invention is not restricted to 'hybrid' sounds.

Furthermore, it is to be noted that the synthesized signal does not need to have the same pitch (fundamental frequency) as the original. In some applications it is required to change the pitch, for example in order to synthesize singing. In order to accomplish this change of fundamental frequency in the synthesized signal, the period locations in the synthesized signal will be placed more closely or more away from each other than the original. This does not otherwise change the synthesis procedure.

Further it is to be noted that the present invention is not restricted to a certain choice of a window. Instead of raised cosine or sine windows other windows can be used such as triangular windows.

CLAIMS:

1. A method of synthesizing a first sound signal based on a second sound signal, the first sound signal having a required first fundamental frequency and the second sound signal having a second fundamental frequency, the method comprising the steps of:

- determining of required pitch bell locations in the time domain of the first sound signal, the pitch bell locations being distanced by one period of the first fundamental frequency,
- providing of pitch bells by windowing the second sound signal on pitch bell locations in the time domain of the second sound signal, the pitch bell locations being distanced by one period of the second fundamental frequency,
- randomly selecting of a pitch bell from the provided pitch bells for each of the required pitch bell locations,
- performing an overlap and add operation on the selected pitch bells for synthesizing the first signal.

2. The method of claim 1, whereby the second sound signal is a hybrid sound comprising a noisy and periodic component.

3. The method of claims 1 or 2, the second sound signal being a voiced fricative sound signal.

4. The method of any one of the preceding claims 1, 2 or 3, the second sound signal being a voiced sound signal and whereby a raised cosine is used for windowing of the second sound signal.

5. The methods of any one of the preceding claims 1, 2 or 3, the second sound signal being an unvoiced sound signal and whereby a sine window is used for windowing of the second sound signal.

6. The method of any one of the preceding claims 1 to 5, the second sound signal having spectrally alike periods, the spectrally alike periods having basically the same information content.

7. The method of any one of the preceding claims 1 to 6, the required first fundamental frequency and the second fundamental frequency being substantially the same.

8. A computer program product, in particular digital storage medium, comprising program means for synthesizing of a first sound signal based on a second sound signal, the first sound signal having a required first fundamental frequency and the second sound signal having a second fundamental frequency, the program means being adapted to perform the steps of:

- determining of required pitch bell locations in the time domain of the first sound signal, the pitch bell locations being distanced by one period of the first fundamental frequency,
- providing of pitch bells by windowing the second sound signal on pitch bell locations in the time domain of the second sound signal, the pitch bell locations being distanced by one period of the second fundamental frequency,
- randomly selecting of a pitch bell from the provided pitch bells for each of the required pitch bell locations,
- performing an overlap and add operation on the selected pitch bells for synthesizing the first signal.

9. A computer system, in particular text-to-speech synthesis system, for synthesizing a first sound signal based on a second sound signal, the first sound signal having a required first fundamental frequency and the second sound signal having a second fundamental frequency, the computer system comprising:

- means for determining of required pitch bell locations in the time domain of the first sound signal, the pitch bell locations being distanced by one period of the first fundamental frequency,
- means for providing of pitch bells by windowing the second sound signal on pitch bell locations in the time domain of the second sound signal, the pitch bell locations being distanced by one period of the second fundamental frequency,

- means for randomly selecting of a pitch bell from the provided pitch bells for each of the required pitch bell locations,
- means for performing an overlap and add operation on the selected pitch bells for synthesizing the first signal.

10. The computer system of claim 9 further comprising means for storing of sound classification data, the means for storing of sound classification data being adapted to store data being indicative of an interval containing the second sound signal within an original sound signal.

11. A synthesizing signal comprising a number of pitch bells which are overlapped and added, each of the pitch bells being randomly selected from a set of pitch bells which are obtained by windowing of an original sound signal on pitch bell locations in the time domain of the second sound signal, the pitch bell locations being distanced by one period of the fundamental frequency.

ABSTRACT:

The present invention relates to a method of synthesizing a first sound signal based on a second sound signal, the first sound signal having a required first fundamental frequency and the second sound signal having a second fundamental frequency, the method comprising the steps of:

- determining of required pitch bell locations in the time domain of the first sound signal, the pitch bell locations being distanced by one period of the first fundamental frequency,
- providing of pitch bells by windowing the second sound signal on pitch bell locations in the time domain of the second sound signal, the pitch bell locations being distanced by one period of the second fundamental frequency,
- randomly selecting of a pitch bell from the provided pitch bells for each of the required pitch bell locations,
- performing an overlap and add operation on the selected pitch bells for synthesizing the first signal.

Fig. 3

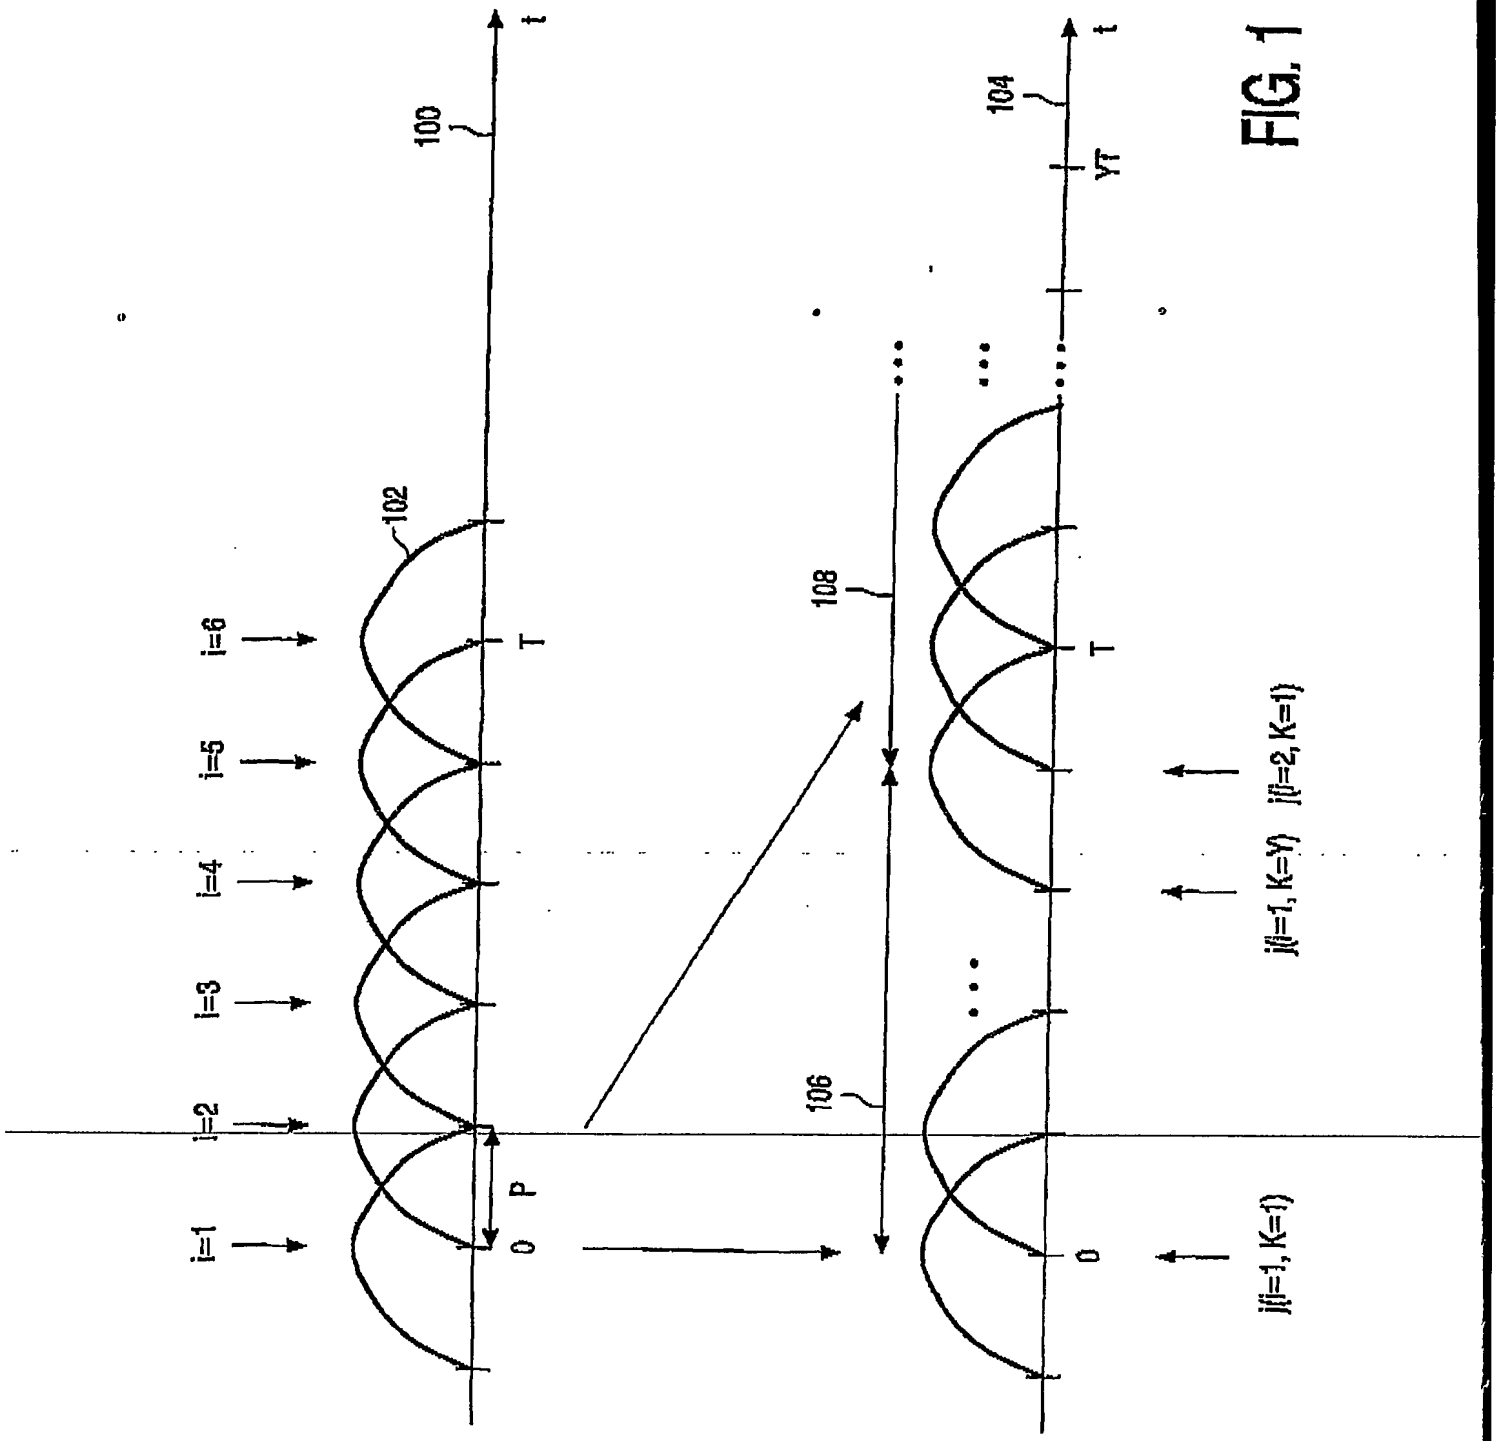

FIG. 1

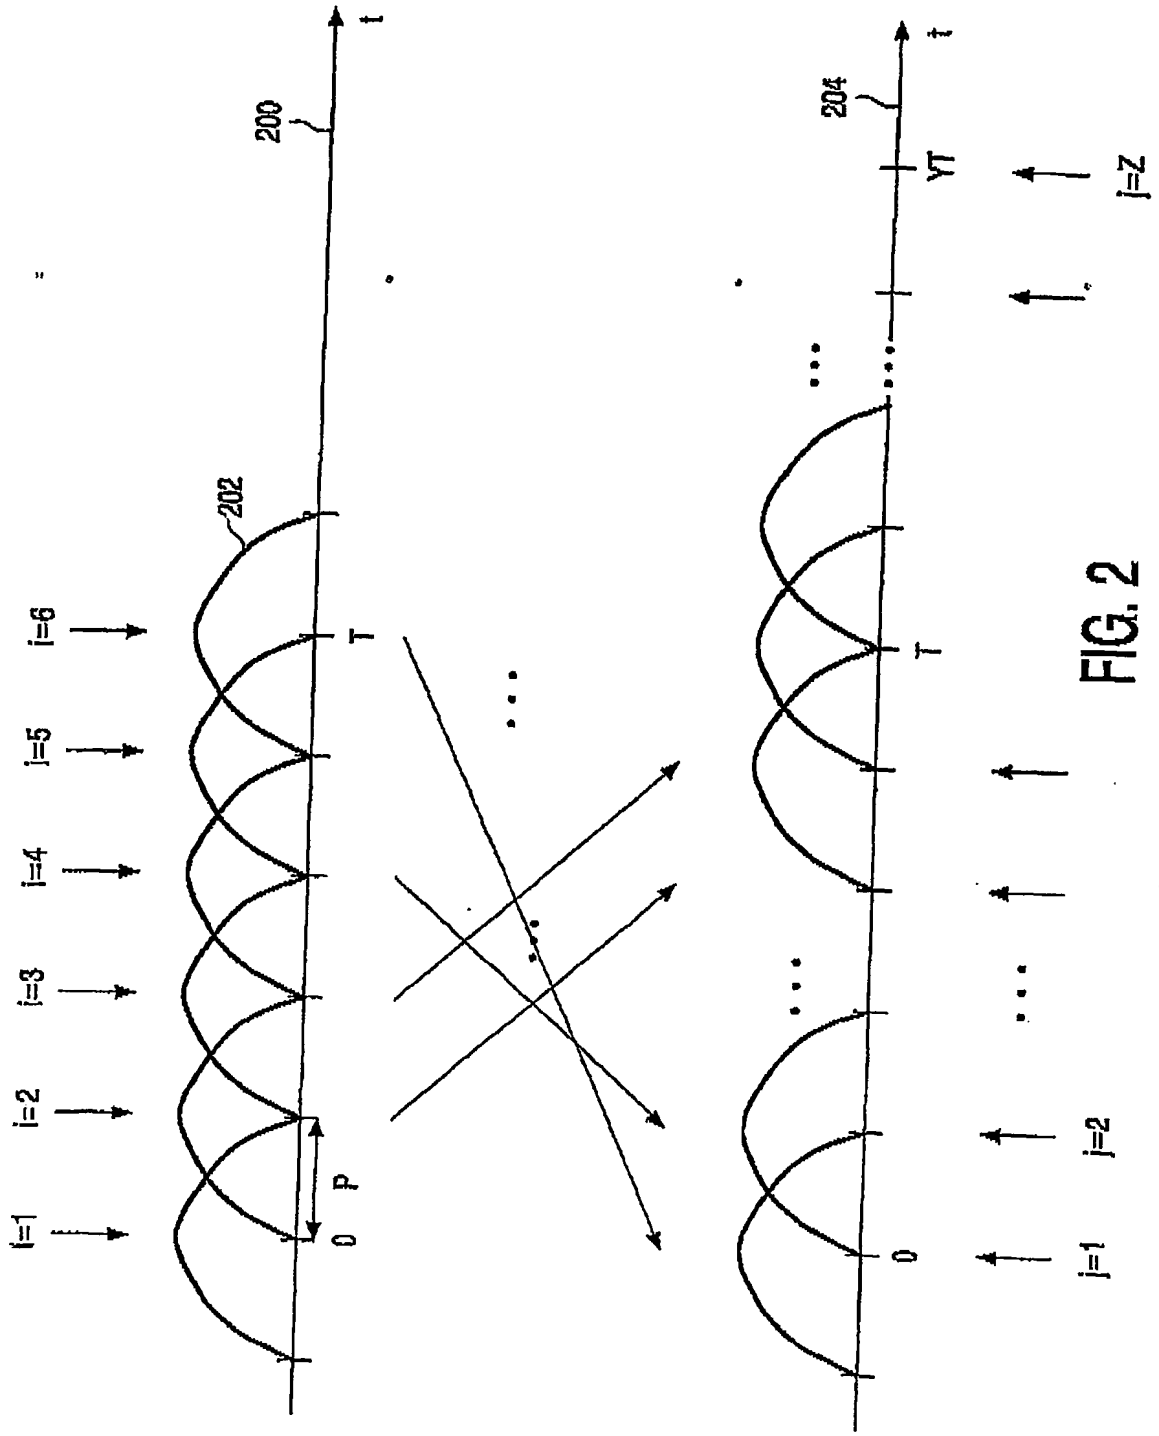

FIG. 2

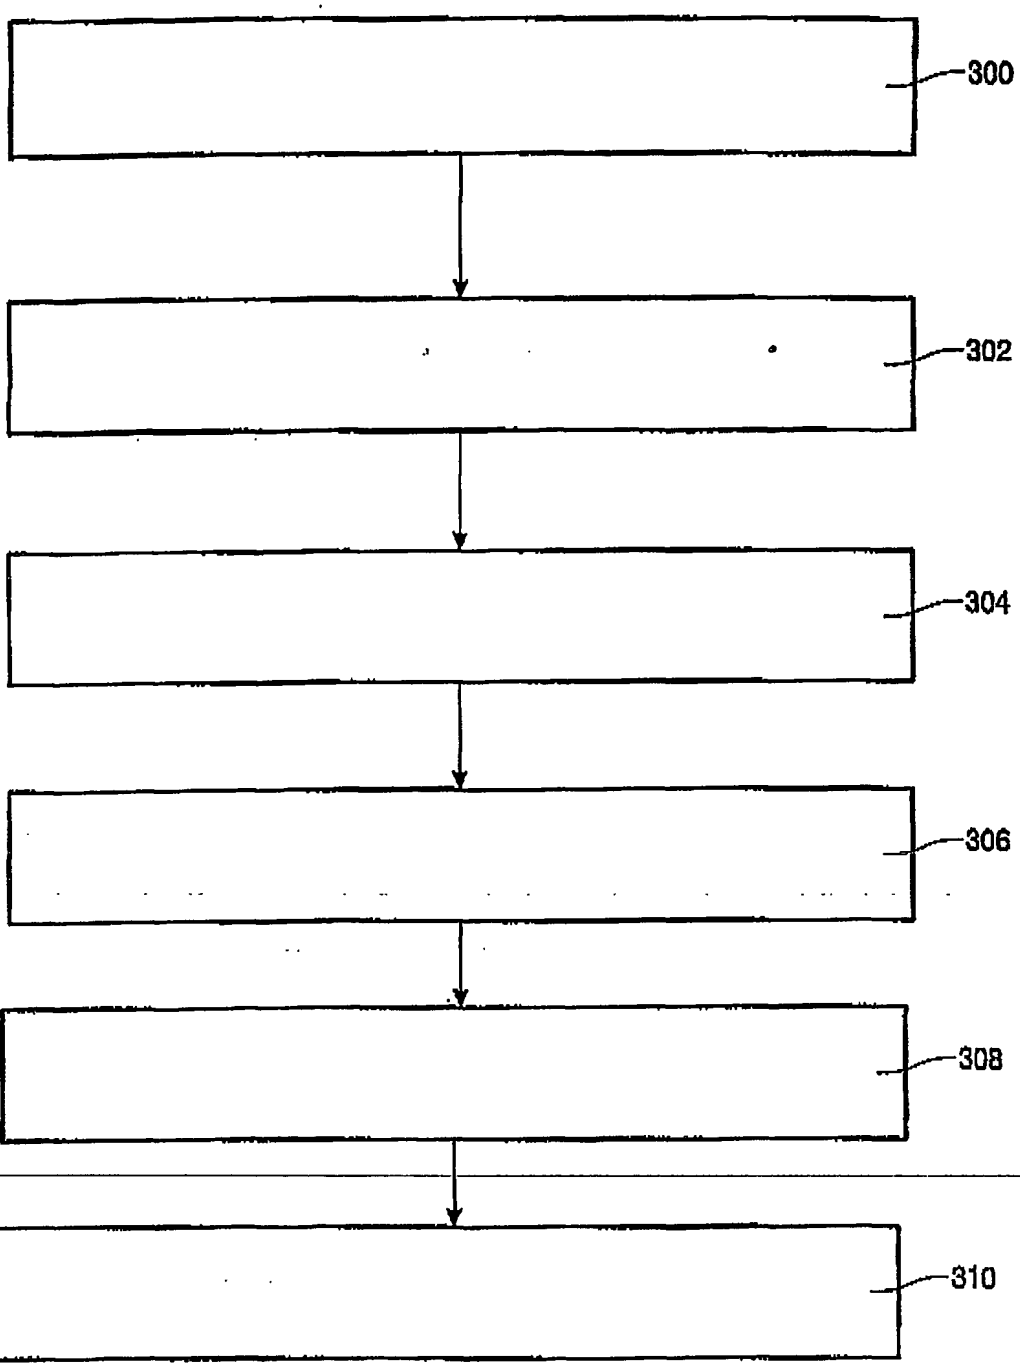

FIG. 3

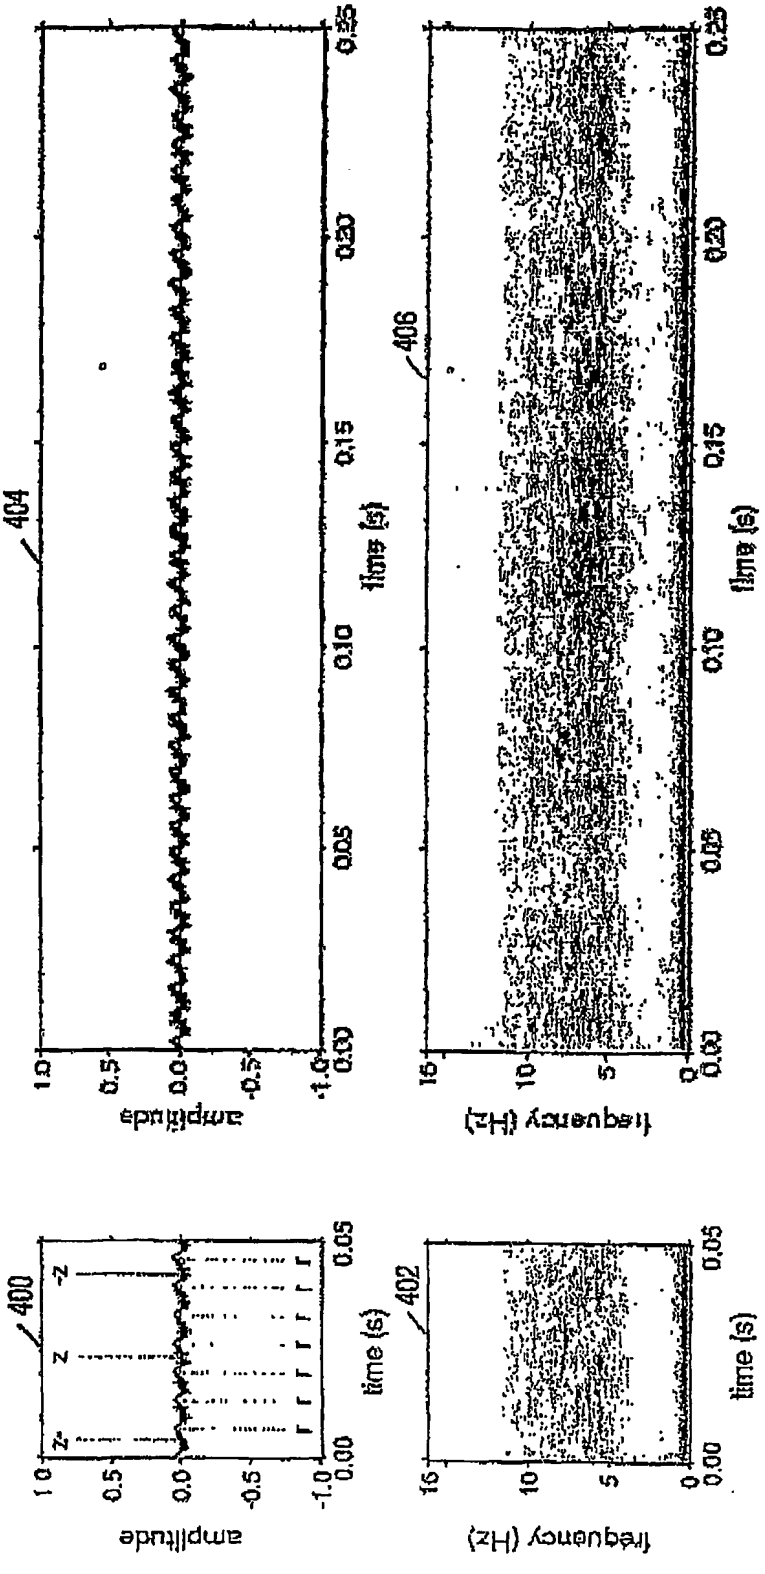

FIG. 4

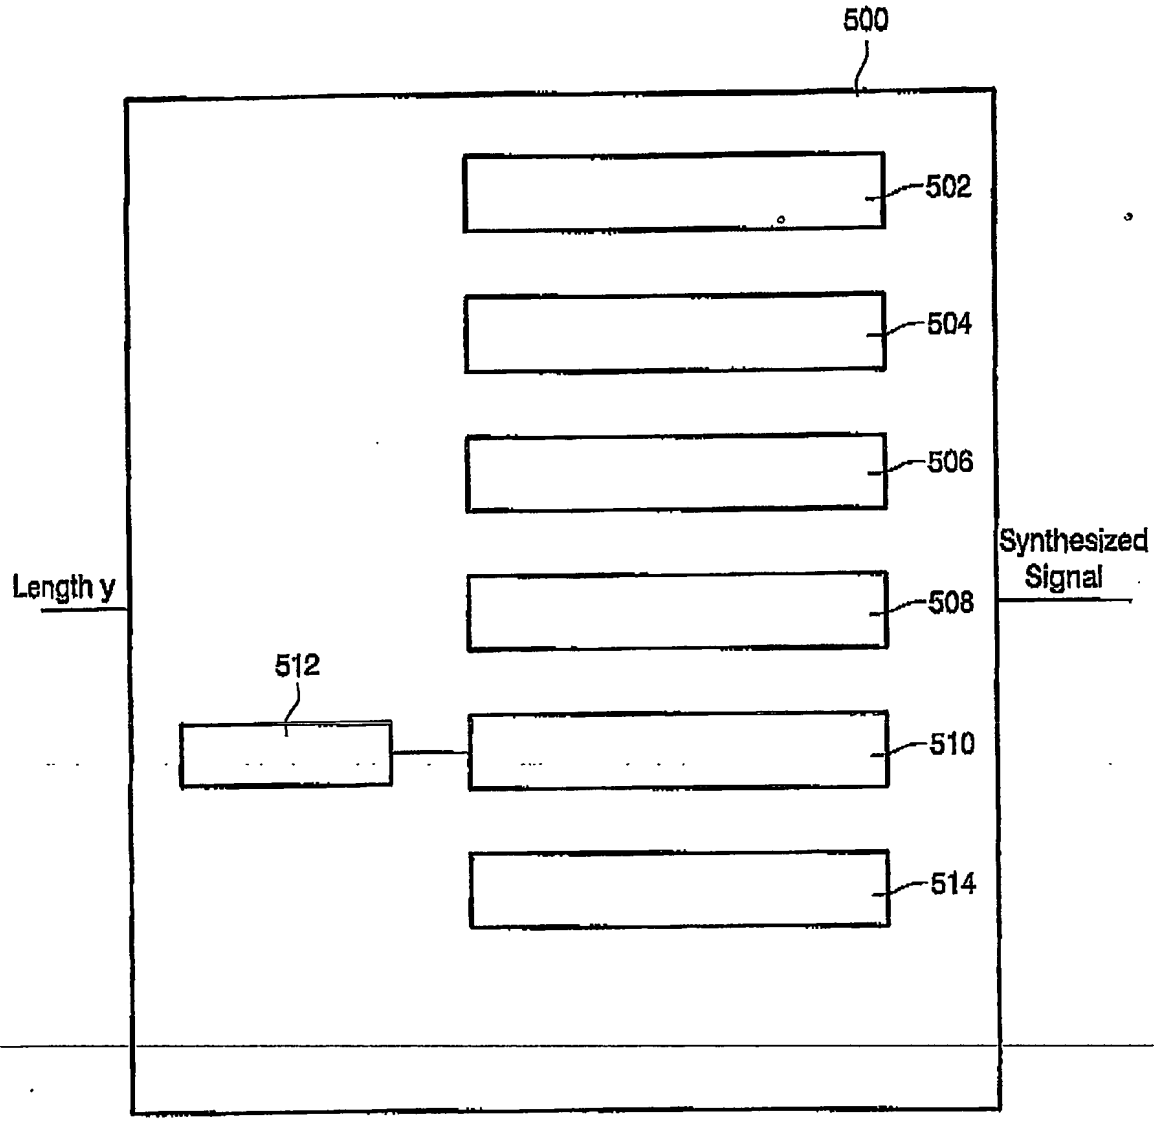

FIG. 5